arXiv:1510.07309v1 [math.PR] 25 Oct 2015

SCALED SUBORDINATORS AND GENERALIZATIONS OF THE INDIAN BUFFET PROCESS

Lancelot F. James, Peter Orbanz and Yee Whye Teh

HKUST, Columbia University and University of Oxford

We study random families of subsets of $\mathbb{N}$ that are similar to exchangeable random partitions, but do not require constituent sets to be disjoint: Each element of $\mathbb{N}$ may be contained in multiple subsets. One class of such objects, known as Indian buffet processes, has become a popular tool in machine learning. Based on an equivalence between Indian buffet and scale-invariant Poisson processes, we identify a random scaling variable whose role is similar to that played in exchangeable partition models by the total mass of a random measure. Analogous to the construction of exchangeable partitions from normalized subordinators, random families of sets can be constructed from randomly scaled subordinators. Coupling to a heavy-tailed scaling variable induces a power law on the number of sets containing the first n elements. Several examples, with properties desirable in applications, are derived explicitly. A relationship to exchangeable partitions is made precise as a correspondence between scaled subordinators and Poisson-Kingman measures, generalizing a result of Arratia, Barbour and Tavaré on scale-invariant processes.

1. Overview and main results. The Indian buffet process, or IBP, of Griffiths and Ghahramani [15] is a distribution on families of sets, encoded as binary matrices. This model and its two- and three-parameter generalizations [40, 38] have received considerable attention in machine learning [14, 42, 41, 32, 7, 17], where applications include dyadic data [28], link prediction [29], time series [13], user preference data [30], and networks [9]. Theoretical results include the work of Broderick, Jordan, and Pitman [8], who generalize exchangeable partition probability functions to families of sets, and of Berti, Crimaldi, Pratelli, and Rigo [4], who establish central limit theorems.

Random families of sets are related to random partitions of $\mathbb{N}$—the former differ from the latter in that each $n \in \mathbb{N}$ may be contained in multiple blocks, or in no block at all. The law of an exchangeable random partition can be parametrized by a random probability measure, via Kingman's representation [see 36]. Similarly, the laws of a large class of random families of sets can be parametrized by random measures, though these are not normalized [40].

The random measure defining the one-parameter IBP of Griffiths and Ghahramani [15] derives from the scale-invariant Poisson process [e.g. 1, 33]. We use this process to relate the IBP to the one-parameter Chinese restaurant process, via a remarkable result of Arratia, Barbour, and Tavaré [2]. Beyond the scale-invariant case, a random scaling variable arises naturally, and we show that Poisson-Kingman partitions [35] correspond to a class of random set families defined by randomly scaled subordinators. For the IBP and CRP, this result specializes to that of Arratia *et al.* [2]. The random measures so defined have a number of

Primary 60G57, 60C05; secondary 60E99, 60G52

Keywords and phrases: Random measures, Indian Buffet process, scale invariant process, Dickman distribution, machine learning, relational models

generic properties, including a "stick-breaking" representation. Substituting a stable for the scale-invariant subordinator, we obtain explicit examples, some with properties reminiscent of the two-parameter Poisson-Dirichlet [37]. If the scaling variable is chosen heavy-tailed, the number of subsets containing the first n elements of $\mathbb{N}$ follows a power law. A different but related construction, based on scaling by a variable derived by Bertoin, Fujita, Roynette, and Yor [5], yields simultanuous power laws on the number and the sizes of subsets.

We begin with a summary of the main results, including a brief review of definitions and preliminaries, in the remainder of this. Technical details and specific examples follow in Sections 2–6. All proofs are collected in the appendix.

1.1. *Background.* A **latent feature model** generates a random binary $n \times \infty$ matrix $\mathbf{Z}$ and a random sequence $\mathbf{U} = (U_1, U_2, \ldots)$ with entries in some space $\mathbb{U}$. Suppose data $X_1, \ldots, X_n$ is observed, say in $\mathbb{R}^d$, where each observation X_i is a list of measurements describing an object i. Given a suitable family of distributions P_ϕ on the sample space, a latent feature model explains such data as

$$X_{ij} \sim P_{\phi_{ij}} \qquad \text{where } \phi_{ij} := (\mathbf{Z}_{i1}U_1, \mathbf{Z}_{i2}U_2, \ldots) \; . \tag{1.1}$$

The model assumes an unobserved set of properties, called **features**; each object i may posses feature k ($\mathbf{Z}_{ik} = 1$), or not possess it ($\mathbf{Z}_{ik} = 0$). The variable $\mathbf{Z}_{ik}$ acts as a switch that turns the effect of parameter U_j on X_{ij} on or off. Throughout, we always require

$$\textstyle\sum_k \mathbf{Z}_{ik} < \infty \qquad \text{for all } i = 1, \ldots, n \; . \tag{1.2}$$

Various forms of this model are widely used in both machine learning and statistics. Here, $\mathbf{Z}$ and $\mathbf{U}$ constitute the model parameters; since both are random variables, the model is Bayesian. Since both have an infinite number of columns, the model is nonparametric; a parametric model would be obtained by restricting $\mathbf{Z}$ and $\mathbf{U}$ to a fixed, finite number of columns. Typically, however, only a finite (if unbounded) number of parameters should explain any single observation, which is ensured by the constraint (1.2). Perhaps the most widely used example of a latent feature model assumes the parameters U_j are vectors of scalar effects that combine linearly, and the only additional randomness are additive, independent noise contribution ε_{ij}, in which case (1.1) takes the form $X_{ij} = \sum_k \mathbf{Z}_{ik}U_{jk} + \varepsilon_{ij}$. This model and its variants are encountered in statistics as linear factor analyzers, and in machine learning and as collaborative filtering methods; see e.g. [16].

In more probabilistic terms, latent features models are closely related to random partitions of the integers: Encode a partition of $\mathbb{N}$ into disjoint blocks as a binary matrix $\mathbf{Z}$, where $\mathbf{Z}_{ik} = 1$ if i is in block k of the partition, and $\mathbf{Z}_{ik} = 0$ otherwise. Since blocks are disjoint, $\sum_k \mathbf{Z}_{ik} = 1$ for all i. A latent feature model relaxes this row sum constraint to the finiteness constraint (1.2). For example, for the subset $\{1, \ldots, 6\}$, the two matrices

$$\begin{array}{c} \\ \text{object \#} \end{array} \begin{array}{c} \text{block \#} \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{array} \qquad \begin{array}{c} \\ \text{object \#} \end{array} \begin{array}{c} \text{block \#} \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

respectively represent the partition and family of sets

$$(\{1,2,4\},\{3\},\{5\},\{6\}) \qquad \text{and} \qquad (\{1,2,3,4\},\{3\},\{4,5\},\{3,6\}) \; . \tag{1.3}$$

The random matrix $\mathbf{Z}$ in a latent feature model thus encodes a random family of sets. In the following, we consider two types of representations for the law of $\mathbf{Z}$:

- **Urn schemes** generate each row of $\mathbf{Z}$ conditionally on the previous ones. Following the interpretation of $\mathbf{Z}$ as a generalization of a random partition to a family of sets, these urn schemes can be regarded as a generalization of the Blackwell-MacQueen urn scheme.
- **Hierarchical representations** first generate a random discrete measure

$$\xi(\bullet) \stackrel{\text{d}}{=} \sum_{k\in\mathbb{N}} J_k \delta_{U_k}(\bullet) \qquad \text{where } J_1, J_2, \ldots \in [0,1] \text{ a.s.} \tag{1.4}$$

  on $\mathbb{U}$. The matrix $\mathbf{Z}$ is then generated column-wise as

$$\mathbf{Z}_{1k}, \ldots, \mathbf{Z}_{nk} | (J_k) \sim_{\text{iid}} \text{Bernoulli}(J_k) \; . \tag{1.5}$$

The hierarchical representation generates matrices whose row vectors form an exchangeable sequence. That is not typically true for urn schemes, since each row is generated conditionally on the previous one. The encoding of family of sets by $\mathbf{Z}$ is not unique, however, and we call two matrices **equivalent** if both encode the same family of sets, and denote the equivalence class of $\mathbf{Z}$ as $[\mathbf{Z}]$. We call $\mathbf{Z}_1$ and $\mathbf{Z}_2$ **equivalently distributed** if $[\mathbf{Z}_1] \stackrel{\text{d}}{=} [\mathbf{Z}_2]$.

The urn scheme representation was first used by Griffiths and Ghahramani [15], who introduced a latent feature model which they called the **Indian buffet process** (IBP). To specify urn schemes, we frequently rely on the quantities

$$n_k := \sum_{i=1}^{n} \mathbf{Z}_{ik} \qquad \text{and} \qquad K_n := \sum_k \mathbb{I}\{n_k > 0\} \; , \tag{1.6}$$

i.e. in terms of the object/feature interpretation, n_k of the first n observed object possess feature k, and K_n is the total number of distinct feature exhibited by the first n objects. The following version is due to Ghahramani, Griffiths, and Sollich [14]: Choose $c, \theta > 0$.

1. The first row of $\mathbf{Z}$ contains $K_1 \sim \text{Poisson}(c)$ consecutive non-zero entries.
2. In row $n+1$, generate the first K_n entries as $\mathbf{Z}_{n+1,k} \sim \text{Bernoulli}\big(\frac{n_k}{\theta+n}\big)$, then append $\text{Poisson}\big(\frac{c\theta}{\theta+n}\big)$ consecutive non-zero entries.

By fixing $\theta = 1$, one obtains the original, one-parameter IBP in [15].

The hierarchical representation using a random measure was introduced as an alternative construction of the IBP by Thibaux and Jordan [40], who showed that the matrix $\mathbf{Z}$ generated by (1.5) is equivalently distributed to an IBP(c,θ) matrix if the sequence (J_k) is generated as the jumps of a subordinator with Lévy density

$$\lambda(s) = \frac{c}{\theta} s^{-1} (1-s)^{\theta-1} \; . \tag{1.7}$$

A random measure $\xi = \sum_k J_k \delta_{U_k}$ is called **homogeneous** if the atom locations $U_1, U_2, \ldots$ are i.i.d. variables and independent of the sequence (J_k). If the weight sequence (J_k) of a homogeneous random measure is generated by a subordinator, then ξ is in particular a **completely random measure** (CRM) in the sense of Kingman [23]. Completely random measures with Lévy density (1.8) were introduced in this form by Hjort [18], who called them **beta processes**. They also appear, up to a transformation $s \mapsto -\log(1-s)$, in [12]. Teh and Görur [38] generalize (1.8), with an additional parameter $\alpha > -\theta$, as

$$\lambda(s)ds = \frac{c}{B(\alpha+\theta, 1-\alpha)} s^{-\alpha-1}(1-s)^{\theta+\alpha-1} ds \tag{1.8}$$

where B is the beta function. The resulting CRM is called the **stable-beta process**. Via (1.5), the stable-beta process defines a three-parameter generalization $\mathrm{IBP}(c, \theta, \alpha)$ of the Indian buffet process. For $\alpha > 0$, the weights of a stable-beta CRM exhibit a power law; hence, the column sums of the resulting random matrix also have a power law distribution. Note (1.8) is of the form $\lambda(s) = \frac{1}{s} f(s)$, where f is the density of the structural distribution of a two-parameter Poisson-Dirichlet process [36, §2.3]. This identity implies the stick-breaking representation of [32] and [6].

1.2. *Desiderata.* In principle, a distribution for $\mathbf{Z}$ can be specified by choosing any suitable random measure ξ and applying (1.5). An urn scheme can then be obtained by integrating out the random measure ξ and conditioning each row of $\mathbf{Z}$ on the previous rows. For a model to be useful, however, it should satisfy a number of properties:

- The hierarchical representation should exist, i.e. the rows of $\mathbf{Z}$ should be rendered conditionally independent by some random measure.
- The random measure ξ should be tractable, i.e. it should be possible to simulate its weights and atoms. The beta and stable-beta process above can both be simulated using "stick-breaking constructions".
- The urn scheme should be tractable, i.e. it should be possible to sample from conditionals of the form $\mathcal{L}(\mathbf{Z}_{n+1}|\mathbf{Z}_1, \ldots, \mathbf{Z}_n)$, where $\mathbf{Z}_i$ denotes the ith row of $\mathbf{Z}$. Note the Griffiths-Ghahramani urn scheme requires only Bernoulli and Poisson variables.
- Ideally, the conditional $\mathcal{L}(\xi|\mathbf{Z}_1, \ldots, \mathbf{Z}_n)$ should also be available. It can be used to derive the urn scheme from (1.5), and may also be of interest in its own right as the posterior distribution of ξ in the sense of Bayesian statistics.

Although most models will only admit either a hierarchical representation or a tractable urn scheme, the analogy to random partitions suggests some distinguished cases may exist—such as the one- and two-parameter Poisson-Dirichlet in the partition case—for which several or all above properties hold. At present, the IBP family seems to be the only known class of models for which that is the case; our objective in the following is to identify others.

1.3. *Main results.* For use in the hierarchical construction (1.5), a random measure must almost surely satisfy

$$J_1, J_2, \ldots \in [0,1] \qquad \text{and} \qquad \xi(\mathbb{U}) < \infty\,, \tag{1.9}$$

where the second condition ensures (1.2). We call a discrete measure **unitary** if it satisfies (1.9).

In the following, we consider a class $\mathrm{GD}(\lambda, P^\circ)$ of unitary random measures parametrized by a continuous Lévy density λ and a probability measure P°, both defined on $\mathbb{R}_+$. We refer to these measures as **generalized Dickman measures**, for reasons explained in Section 3. Let $\Delta_1 > \Delta_2 > \ldots$ be the jumps of subordinator with Lévy density λ, ordered by decreasing size, and define a random measure ξ as

$$\xi(\bullet) := \textstyle\sum_{k\in\mathbb{N}} J_k \delta_{U_k}(\bullet) \qquad \text{where } J_k := \frac{\Delta_{k+1}}{\Delta_1} , \tag{1.10}$$

so $J_k \in [0,1]$ almost surely. We denote the law of ξ generically by $\mathrm{GD}(\lambda)$.

The variables $\Delta_2, \Delta_2, \ldots$, with the largest jump Δ_1 removed, are conditionally independent given Δ_1, so $\Delta_{2:\infty}|\Delta_1$ is a subordinator. Conditioning on $\Delta_1 = a$ does not change the the Lévy density on $[0,a]$, but truncates it at a. The conditional random measure

$$\xi_a := \xi|(\Delta_1 = a) \tag{1.11}$$

is hence unitary and completely random, with Lévy density

$$\lambda_a(s) = a\lambda(as)\mathbb{I}\{s \le 1\} . \tag{1.12}$$

A natural further step is to randomize a: Choose any non-negative random variable $\Delta^\circ \sim P^\circ$ and define the unitary random measure

$$\xi^\circ \;:=\; \xi_{\Delta^\circ} \;=\; \xi|(\Delta = \Delta^\circ) . \tag{1.13}$$

The law of ξ° is denoted $\mathrm{GD}(\lambda, P^\circ)$. Clearly, choosing Δ° as Δ_1 recovers $\xi^\circ \sim \mathrm{GD}(\lambda)$.

Scaling subordinators yields a natural stick-breaking representation: Since the weights of a random measure $\xi^\circ \sim \mathrm{GD}(\lambda, P^\circ)$ can be expanded as

$$J_k = \frac{\Delta_{k+1}}{\Delta^\circ} = \frac{\Delta_{k+1}}{\Delta_k}\frac{\Delta_k}{\Delta_{k-1}}\cdots\frac{\Delta_2}{\Delta^\circ} , \tag{1.14}$$

ξ° can be represented as

$$\xi^\circ \stackrel{\mathrm{d}}{=} \sum_{k=1}^{\infty}\Big(\prod_{j=1}^{k} R_k\Big)\delta_{U_k} \qquad \text{where } R_k := \frac{\Delta_{k+1}}{\Delta_k} \text{ and } R_1 := \frac{\Delta_2}{\Delta^\circ} . \tag{1.15}$$

Following Kingman [24], the density of the largest jump of a subordinator can always be obtained explicitly as

$$f_{\Delta_1}(s) = \lambda(s)e^{-\Lambda(s)} \qquad \text{where} \qquad \Lambda(s) := \int_s^\infty \lambda(u)du . \tag{1.16}$$

Whenever the function Λ has a computationally tractable inverse, the representation above can be made explicit:

THEOREM 1.1. *For any random measure $\xi^\circ \in \mathrm{GD}(\lambda, P^\circ)$, define the right-continuous inverse $\Lambda^-(s) := \inf\{t|\Lambda(t) \ge s\}$ of Λ in (1.16), and let $E_1, E_2, \ldots \sim_{\mathrm{iid}} \mathrm{Exponential}(1)$. Then for any $a > 0$,*

$$(\Delta_2, \Delta_3, \ldots)\big|\Delta^\circ = a \quad \stackrel{\mathrm{d}}{=} \quad (M_2^{(a)}, M_3^{(a)}, \ldots) \tag{1.17}$$

where $M_k^{(a)} := \Lambda^-\big(\Lambda(a) + E_2 + \ldots + E_k\big)$ for every $k \ge 2$.

One can hence simulate ξ° by generating $\Delta^\circ \sim P^\circ$ and

$$\xi^\circ|(\Delta^\circ = a) \stackrel{\mathrm{d}}{=} \sum_{k\in\mathbb{N}} \Big(\prod_{j=1}^{k} \frac{M_{k+1}^{(a)}}{M_k^{(a)}} \Big) \qquad \text{where} \qquad M_1^{(a)} := a \; . \tag{1.18}$$

In the stable case discussed in Section 3, Λ is indeed invertible, and the random variables $M_k^{(a)}$ have explicit representations in terms of $\Lambda^- = \Lambda^{-1}$. Even if Λ is not invertible, it is typically still possible to sample $M_k^{(a)}$.

In terms of the desiderata listed above, a matrix $\mathbf{Z}$ generated from a $\mathrm{GD}(\lambda, P^\circ)$ measure with a sufficiently regular function Λ thus has a hierarchical representation by a tractable random measure; that leaves the conditional $\mathcal{L}(\xi|\mathbf{Z}_1, \ldots, \mathbf{Z}_n)$ and the urn scheme. General results on conditional distributions of ξ and $\mathbf{Z}_{n+1}$ are given in Section 2. We mention one elementary (but apparently overlooked) fact that may be of independent interest in modeling and simulation problems: If $\mu = \sum_k W_k \delta_{U_k}$ is a homogeneous CRM with Lévy density λ, and h is a positive function, the random measure

$$\mu_h := \sum_k B_k W_k \delta_{U_k} \qquad \text{where} \qquad B_k|\mu \sim \mathrm{Bernoulli}(h(B_k)) \tag{1.19}$$

is distributed as a homogeneous CRM with Lévy density $h\lambda$. Thus, tilting by h can be simulated by thinning; see Lemma 2.4.

To obtain tractable urn schemes, one has to consider specific models, and a case of particular interest are Lévy densities of the form

$$\lambda(s) = cs^{-1-\alpha} \qquad \text{for } \alpha \geq 0 \text{ and } c > 0 \; , \tag{1.20}$$

By construction, the measure ξ° is not completely random, since scaling by Δ_1 makes the variables J_k dependent—unless one chooses λ as in (1.20) and $\alpha = 0$. Since (1.20) implies

$$\lambda_a(s) = ca^{-\alpha} s^{-1-\alpha} \; , \tag{1.21}$$

ξ_a does not depend on a for $\alpha = 0$, and ξ° is hence a CRM. If this random measure ξ° is substituted into (1.5), one obtains precisely the one-parameter IBP(c) distribution of Griffiths and Ghahramani [15]. It is interesting to note that the Lévy density is, in this case, the intensity of a scale-invariant Poisson process; this relationship is discussed in more detail in Section 3.1.

For $\alpha > 0$, (1.20) describes a stable subordinator. The one-parameter IBP can hence be regarded as a limiting case of the stable for $\alpha \searrow 0$. Stick-breaking and conditional probabilities for the stable case are covered in Section 3. One implication of these results is that the stable case admits a simple urn scheme: If $\xi^\circ \sim \mathrm{GD}(cs^{-\alpha-1}, P^\circ)$ for $\alpha > 0$, and ξ° is substituted for ξ in (1.5), the resulting matrix $\mathbf{Z}$ can equivalently be generated as follows.

1. The first row of $\mathbf{Z}$ contains $\mathrm{Poisson}(\frac{(\Delta^\circ)^{-\alpha}\alpha}{1-\alpha})$ consecutive non-zero entries.
2. In the $(n+1)$st row, each of the first K_n entries is non-zero with probability $\frac{n_k - \alpha}{n+1-\alpha}$. Additionally, $\mathrm{Poisson}(C_n)$ consecutive non-zero entries are appended.

The random variable C_n is a function of a sample from the conditional distribution of Δ° given the previous first n rows of $\mathbf{Z}$, and defined in detail in Corollary 3.3.

Since latent feature models are related to exchangeable partitions, and both are generated by a class of random measures, it is obvious to ask how these classes of random measures are related. Simply normalizing generalized Dickman measures does not generally seem to yield interesting random probability measures. A different picture emerges, however, if one first conditions on a total mass $\xi^\circ(\mathbb{U}) = t$ of $t \leq 1$:

THEOREM 1.2. *Let $\xi^\circ \sim \mathrm{GD}(\lambda, P^\circ)$, for any choice of λ and P°. Condition ξ° on its total mass and normalize, defining the measure*

$$\overline{\eta} := \frac{\xi^\circ | (\xi^\circ(\mathbb{U}) = t)}{t} \,. \tag{1.22}$$

Whenever $t \in (0, 1]$, the random probability measure η is a Poisson-Kingman measure. Conversely, every Poisson-Kingman measure can be obtained in this manner.

One implication is that conditioning on a total mass of at most 1 can drastically simplify the properties of a random measure. This surprising fact was observed by Arratia, Barbour, and Tavaré [2] for the scale-invariant Poisson process, and is arguably implicit in the work of Perman [34]. Our result specializes to that of Arratia, Barbour, and Tavaré [2, Theorem 3.1] for $\lambda(s) = cs^{-1}$, although their proof does not seem to generalize, and our result is obtained in a very different manner. The theorem also shows the random measures $\overline{\eta}$ in (4.4) can be regarded as a refinement of Poisson-Kingman measures, obtained by including the additional mixing variable Δ°. A more detailed statement, including the distributions of the resulting measures, is given in Theorem 4.1.

REMARK 1.3. We note two aspects of possible interest to applications: (i) Given a random partition model with desirable properties, Theorem 1.2 identifies the corresponding feature model; Section 4.3 provides examples. (ii) An explicit coupling between between the matrix $\mathbf{Z}$ and an exchangeable partition can be constructed, by parametrizing both by a single random measure. One can hence specify models for problems where part of the data constitutes a latent feature problem and another part a clustering problem. Problems of this type are known in machine learning as multi-task learning problems, see e.g. [27, 19, 31].

Finally, we consider the distribution of certain functions of $\mathbf{Z}$: The basic statistics of interest in feature modeling problems are typically the row sums (number of features of one object), column sums (number of objects exhibiting a feature), and the total number of features K_n as defined in (1.6). Of particular interest in applications are models describing heavy-tailed phenomena, where one or more of the above statistics exhibit a power law.

Column sums following a power law can be obtained relatively easily, by choosing ξ such that the weight sequence (J_k) is heavy-tailed, as is the case for the stable-beta process (1.8) and for several models in Section 3. We do not consider power laws on the row sums within a single matrix, since—as pointed out in [38] and [7]—this comes at the price of sacrificing the representation (1.5): The row sums $\Sigma_1^{\mathrm{R}}, \Sigma_2^{\mathrm{R}}, \ldots$ are conditionally i.i.d. given ξ, and

$$\mathbb{P}\big(|\Sigma_i^{\mathrm{R}} - \xi(\mathbb{U})| > \varepsilon \,\big|\, \xi\big) \leq \exp\Big(-\frac{\varepsilon^2}{2\xi(\mathbb{U}) - \varepsilon/3}\Big) \qquad \text{for all } \varepsilon > 0 \,, \tag{1.23}$$

by standard Bernoulli tail bounds [see 20]. Within a matrix $\mathbf{Z}$ generated by (1.5), the distribution of row sums is hence not heavy-tailed.

In Section 5, we consider models which result in K_n being marginally heavy-tailed—that is, if $\mathbf{Z}^{(1)}, \mathbf{Z}^{(1)}, \ldots$ are i.i.d. realizations of $\mathbf{Z}$, each with n rows, $K_n(\mathbf{Z}^{(1)}), K_n(\mathbf{Z}^{(1)}), \ldots$ empirically exhibit a power law. That is not the case for the IBP, including the stable-beta case. One way to achieve this behavior is to use a $\mathrm{GD}(\lambda, P^\circ)$ model and choose P° such that $1/\Delta^\circ$ is heavy-tailed. Additionally choosing λ as stable, for example, yields a model with power law for the column sums and a marginal power law for K_n.

There is another, perhaps less obvious construction: If λ is of the form (1.21), the conditional measure $\xi^\circ|(\Delta^\circ = a)$ has Lévy density (1.21). In the scale invariant case $\alpha = 0$, ξ_a does not depend on a. In the stable case $\alpha > 0$, it does, but a only acts as a scaling factor on the Lévy density, and one can hence equivalently sample the weights (J_k) of ξ° as

$$(J_k) \sim \mathrm{Subordinator}(\zeta\lambda) \;, \tag{1.24}$$

where $\zeta := (\Delta^\circ)^{-\alpha}$. Although this representation coincides with the GD representation in the scale-invariant and the stable case, it clearly does not for general choice of λ. In Section 5, we show that there is a specific choice of the variable ζ for which (1.24) nonetheless yields a simple urn scheme; moreover, the distributions of K_n, for $n \in \mathbb{N}$, can be obtained explicitly. For a stable-beta density, for example, we obtain the following: Choose parameters $\sigma \in (0,1)$ and $\alpha, \theta > 0$.

1. The first row contains $\mathrm{Poisson}(\phi_1\zeta)$ consecutive non-zero entries, where $\zeta := G/\beta$, for $G \sim \mathrm{Gamma}(1-\sigma, 1)$ and $\beta \sim \mathrm{Beta}(\sigma, 1)$.
2. In the $(n+1)$st row, each of the first K_n entries is non-zero with with probability $\frac{n_k - \alpha}{\theta + n}$. Additionally, append $H_{n+1}|K_n \sim \mathrm{Poisson}(G_{n+1}F_{n+1})$ consecutive non-zero entries, where $G_{n+1} \sim \mathrm{Gamma}(1-\sigma+K_n)$, and the random variable F_{n+1} is a function of K_n, the parameters, and an independent uniform variable.

The variables ζ and F_n, and the coefficients ϕ_k, are trivial to evaluate; all are described in detail in Section 5, where we also study the distribution of K_n, and note certain parallels between (5.3) and the two-parameter Poisson-Dirichlet distribution.

Since the row sums are not generally Poisson if ξ is not completely random, it stands to reason to ask for a Poisson approximation. In Section 6, we consider a simple total variation bound on the approximation error, and show how, as a consequence of Theorem 1.2, the small weights of unitary random measure can be related to those of a random probability measure. The latter generalizes a result of Arratia, Barbour, and Tavaré [2].

2. Conditioning on a set of rows. Results in this section are largely auxiliary, but requisite for the ensuing discussion. Suppose the first n rows $\mathbf{Z}_1, \ldots, \mathbf{Z}_n$ of a matrix $\mathbf{Z}$ are generated, using either an urn scheme, or a hierarchical representation with random measure ξ. We are concerned with two types of conditional distributions, the laws of the conditional variables $\xi|\mathbf{Z}_1, \ldots, \mathbf{Z}_n$ and $\mathbf{Z}_{n+1,}|\mathbf{Z}_1, \ldots, \mathbf{Z}_n$. In the model (1.1), $\mathbf{Z}$ and $\mathbf{U}$ are used to explain observational data X. In the following, however, it can be useful to interpret the first n rows of $\mathbf{Z}$ itself as n observations. In the terminology of Bayesian statistics, $\mathcal{L}(\xi|\mathbf{Z}_1, \ldots, \mathbf{Z}_n)$ is then the posterior distribution of ξ. Machine learning algorithms, which represent the matrix $\mathbf{Z}$ as a "latent" variable, require a tractable representation of this posterior. We use it in

the following to compute the conditional distribution $\mathcal{L}(\mathbf{Z}_{n+1,}|\mathbf{Z}_1,\ldots,\mathbf{Z}_n)$ of the $(n+1)$st row in an urn scheme.

2.1. *Bernoulli process.* Following Thibaux and Jordan [40], we use the following joint encoding for the variables $\mathbf{Z}$ and $\mathbf{U}$, which is particularly useful for conditioning: Given a (fixed) unitary measure $m = \sum_k w_k \delta_{u_k}$ on $\mathbb{U}$, define a random measure

$$\Pi(\bullet) := \sum_{k\in\mathbb{N}} Z_k \delta_{u_k}(\bullet) \qquad \text{with } Z_k \sim \text{Bernoulli}(w_k) \text{ independently.} \tag{2.1}$$

Π is called a **Bernoulli process** with parameter m in [40], and denoted $\text{BeP}(m)$ below. Let $\mu = \sum_k W_k \delta_{U_k}$ be a unitary random measure, and sample

$$\Pi_1,\ldots,\Pi_n|\xi \sim_{\text{iid}} \text{BeP}(\xi)\ . \tag{2.2}$$

The non-vanishing binary weights of each random measure Π_i can then be interpreted as the non-zero entries of the nth row of a matrix $\mathbf{Z}$, whose distribution is clearly equivalent to that in (1.5). Define $\mathbf{U}$ using the atoms of μ as $\mathbf{U} := (U_1, U_2, \ldots)$. Then Π_i provides the values of precisely those atoms that correspond to non-zero entries in row i.

Since (2.1) thins a point process, it is not hard to believe—and indeed at times tacitly assumed in the literature—that Π should be Poisson random measure if μ is completely random. That is indeed the case:

PROPOSITION 2.1. *Let $\mu = \sum_k J_k \delta_{U_k}$ be a unitary, homogeneous CRM with $G := \mathcal{L}(U_1)$. Then Π, defined by $\Pi|\mu \sim \text{BeP}(\mu)$, is a Poisson random measure with $\mathbb{E}[\Pi] = \mathbb{E}[\mu(\mathbb{U})]G$.*

2.2. *Conditioning.* Suppose we choose $\mu = \xi^\circ$ in (2.2) and draw a sample consisting of n random measures, $\Pi_{1:n} := (\Pi_1,\ldots,\Pi_n)$. To specify the distribution of $\xi^\circ|\Pi_{1:n}$, we have to keep track of the number of features in the sample, and of the prevalence of each feature. Define n_k and K_n as in (1.6). We say that U_k is **observed** in the sample if $n_k > 0$. Since $T_{\xi^\circ} < \infty$ almost surely, the total number K_n of distinct observed atoms is finite almost surely. It is customary to denote the observed atoms as $U_1^*,\ldots,U_{K_n}^*$. Additionally conditioning on Δ° reduces to the CRM case [26].

PROPOSITION 2.2. *The conditional random measure $\xi_n^* := \xi^\circ|\Pi_{1:n}$ is distributed as*

$$\xi_n^* \stackrel{\text{d}}{=} \xi_n^\circ + \sum_{k=1}^{K} J_{nk}^* \delta_{U_k^*}\ . \tag{2.3}$$

Conditionally on Δ°, the unitary measure ξ_n° is completely random, and $J_{nk} \perp\!\!\!\perp_{\Delta^\circ} \xi_n^\circ$. Let f_{Δ° denote the density of Δ° on $\mathbb{R}_+$, and define

$$c_a(n,n_k) := \int_0^a s^{n_k}(1-s)^{n-n_k}\lambda(as)ds \quad \text{and} \quad \psi_n(a) := a\int_0^1 (1-(1-s)^n)\lambda(as)ds\ . \tag{2.4}$$

The conditional random measure $\xi_n^\circ|(\Delta^\circ = a)$ is given by the Lévy density

$$\lambda_n^a(s) := a(1-s)^n\lambda(as)\mathbb{I}\{s\le 1\}\ , \tag{2.5}$$

and the jumps J^*_{nk} *in* (2.3) *by*

$$\mathbb{P}(J^*_{nk} \in ds | \Delta^\circ = a) = \frac{s^{n_k}(1-s)^{n-n_k}\lambda(as)}{c_a(n,n_k)} ds \ . \tag{2.6}$$

The conditional law of Δ° *is*

$$\mathbb{P}(\Delta^\circ \in da | \Pi_{1:n}) \propto f_{\Delta^\circ}(a) e^{-\psi_n(a)} \Big(\prod_{k=1}^{K_n} c_a(n,n_k) \Big) da \ . \tag{2.7}$$

If the random measure ξ° is marginalized out of the hierarchical model, we obtain the predictive distribution of $\Pi_{n+1}|\Pi_{1:n}$.

PROPOSITION 2.3. *Let* $U^*_1, \ldots, U^*_{K_n}$ *be observed atoms in* $\Pi_{1:n}$*. Under the predictive distribution* $\mathcal{L}(\Pi_{n+1}|\Pi_{1:n})$*, each previously observed atom* U^*_k *has non-zero weight with probability*

$$\mathbb{P}[\mathbf{Z}_{n+1,k} = 1 | \Pi_{1:n}] = \frac{c_a(n+1, n_k+1)}{c_a(n,n_k)} \ . \tag{2.8}$$

Additionally, there are Poisson(q_n) *previously unobserved atoms with non-zero weights, where*

$$q_n := \int_0^1 s(1-s)^n \lambda(s) ds \ , \tag{2.9}$$

and the locations of the newly observed atoms are drawn i.i.d. from G*.*

2.3. *Tilting and thinning.* The results above show conditioning a random measure with Lévy density λ on Bernoulli process observations leads to Lévy densities $\lambda(s)(1-s)^n$. Terms of the form $(1-s)^t$, for some scalar t, also arise in the Lévy densities of the beta stable beta process (1.8). The following thinning lemma shows that such terms, and more generally Lévy densities $\lambda(s)h(s)$ "tilted" by some function h, can be regarded as the outcome of a conditional thinning operation:

LEMMA 2.4. *Let* $\mu = \sum_{k\in\mathbb{N}} J_k \delta_{U_k}$ *be a homogeneous CRM on* $\mathbb{U}$*, with Lévy density* λ *on* $(0,\infty)$*, and let* $h : (0,\infty) \to [0,1]$ *be measurable. If* (B_k) *are conditionally independent binary variables with conditional law* $B_k|J_k \sim$ *Bernoulli*$(h(J_k))$*, then*

$$\mu_h := \sum_{k\in\mathbb{N}} B_k J_k \delta_{U_k} \tag{2.10}$$

is again a homogeneous CRM, with Lévy density $h \cdot \lambda$*. If in particular* $\int_0^\infty h(s)\lambda(s)ds = \infty$*, then* $\mu_h(\mathbb{U}) = \infty$ *almost surely.*

If a procedure for sampling jumps from λ—such as a stick-breaking representation—is available, one can hence sample $\lambda(s)(1-s)^n$ using $B_k \sim \text{Bernoulli}((1-J_k)^n)$. In the GD$(\lambda, P^\circ)$ case, this is applicable to $\xi^\circ_n|(\Delta^\circ = a)$.

3. The scale-invariant and the stable case. We have already noted that both the basic IBP and the CRP are inherently related to the subordinator with Lévy density $\lambda(s) = \theta s^{-1}$ for some $\theta > 0$. In this section, we consider more generally the case

$$\lambda(s) = cs^{-1-\alpha} \qquad \text{for } \alpha, c > 0 \,, \tag{3.1}$$

i.e. $\Delta_1, \Delta_2, \ldots$ are the jumps of a scale-invariant Poisson process (if $\alpha = 0$), or of a stable subordinator (if $\alpha > 0$).

3.1. *Dickman distributions.* The scale-invariant Poisson process derives its name from the fact that a Poisson process $\{X_1, X_2, \ldots\}$ on $\mathbb{R}_+$ satisfies the scale-invariance

$$\{X_1, X_2, \ldots\} \stackrel{\text{d}}{=} \{bX_1, bX_2, \ldots\} \stackrel{\text{d}}{=} \{1/X_1, 1/X_2, \ldots\} \qquad \text{for every } b > 0 \tag{3.2}$$

if and only if $\lambda = cs^{-1}$ for some $c > 0$. The scale invariance gives rise to unique properties; see [10, Chapter 12.8] or [1] for more details.

Let $T := \sum_k X_k$ be the sum of locations of a Poisson process $\{X_1, X_2, \ldots\}$ on $\mathbb{R}_+$ with an arbitrary Lévy density λ. Under mild regularity conditions, T has a probability density g; if so, this density satisfies the integral equation

$$tg(t) = \int_0^\infty g(t-s)s\lambda(s)ds \,, \tag{3.3}$$

see e.g. [35]. Now consider the scale-invariant Poisson process. If this process is defined on all of $\mathbb{R}_+$, its total mass is infinite with positive probability. We hence restrict the process to the unit interval, which ensures $T < \infty$ almost surely. Then $\lambda(s) = \frac{c}{s}\mathbb{I}\{s \leq 1\}$, and

$$tg(t) = c\int_0^1 g(t-s)ds = c\big(G(t) - G(t-1)\big) \,, \tag{3.4}$$

where G denotes the cdf of g. Differentiating on both sides yields

$$t\frac{dg}{dt}(t) = (c-1)g(t) - cg(t-1) \,. \tag{3.5}$$

Continuous solutions g to (3.5) exist and are uniquely determined up to scaling. For $c = 1$, (3.5) is called **Dickman's equation**, and the solution g uniquely determined by $g(1) = 1$ is the **Dickman function** [e.g. 3, 33]. We are interested only in continuous solutions which are probability densities. These are uniquely determined for each $c > 0$. For each c, an $\mathbb{R}_+$-valued random variable D_c with density g is called a **generalized Dickman** variable. Thus, the total mass T of a subordinator with Lévy density $\lambda(s) = cs^{-1}$ has law $\mathcal{L}(D_c)$.

3.2. *Stick-breaking.* For the one-parameter IBP, and hence for the scale-invariant case $\alpha = 0$, a stick-breaking representation was proposed by Teh, Görur, and Ghahramani [39]. It is recovered from (1.15) by choosing $R_k \sim \text{Beta}(c, 1)$.

REMARK 3.1. Like the model itself, the representation in [39] can be understood as an outcome of scaling a subordinator: If a point process is scale-invariant, its image under a

logarithm is translation-invariant. If $\{X_1, X_2, \ldots\}$ is scale-invariant Poisson with parameter c, the process $\{-\log X_1, -\log X_2, \ldots\}$ is again Poisson, now with constant rate c, so its arrival times are distributed as $-\log X_k \stackrel{\mathrm{d}}{=} E_1 + \ldots + E_k$ for independent Exponential$(1/c)$ variables E_k [see 1]. If U denotes a Uniform$[0,1]$ variable, $\mathcal{L}(e^{-E_k}) = \mathcal{L}(U^{1/c}) = \mathrm{Beta}(c, 1)$.

Now consider a general GD(λ, P°) model. In the stable case, Theorem 1.1 shows that the coefficients $M_k^{(a)}$ in (1.18) are

$$M_k^{(a)} = \big(\frac{\alpha}{c}(E_1 + \ldots + E_{k-1}) + a^{-\alpha}\big)^{-1/\alpha} . \tag{3.6}$$

Another specific example is the beta process density $\lambda(s) = \alpha s^{-1-\alpha}(1-s)^{\alpha-1}$, for which

$$M_k^{(a)} = \big(\big(\big(\tfrac{1-a}{a}\big)^{\alpha} + E_1 + \ldots + E_{k-1}\big)^{1/\alpha} + 1\big)^{-1} . \tag{3.7}$$

The distribution of the variables R_k now depends on the choice of P°. It is interesting to note that the law Beta$(c, 1)$ in the IBP case above precisely matches that occurring in the stick-breaking construction of the one-parameter Poisson-Dirichlet process. The corresponding laws in the two-parameter Poisson-Dirichlet are $R_k \sim \mathrm{Beta}(\theta + \alpha k, 1)$. Note that, if $\alpha = 0$, the one-parameter case above is obtained for $\theta = c$. Thus, θ and c play the same role when comparing the one- and two-parameter Poisson-Dirichlet to each other, whereas in the three-parameter models considered here, they are distinct. In a stable GD model, jumps of the form Beta$(\theta + \alpha k, 1)$ are obtained for a specific choice of P°:

COROLLARY 3.2. *If* $\xi^\circ \sim \mathrm{GD}(\lambda(s) = cs^{-1-\alpha}, P^\circ)$ *for* $\alpha > 0$, *and* $P^\circ := \mathrm{Gamma}(\frac{\theta+\alpha}{\alpha}, \frac{c}{\alpha})$ *with* $\theta > -\alpha$, *each coefficient* R_k *has law* $\mathrm{Beta}(\theta + \alpha k, 1)$.

3.3. *Conditional distributions.* If λ is scale-invariant, Proposition 2.2 shows the random measure component ξ_n° of $\xi^\circ|\Pi_{1:n}$, additionally conditioned on $\Delta^\circ = a$, has Lévy density

$$\lambda_{n,a}(s) = cs^{-1}(1-s)^n , \tag{3.8}$$

and is hence a beta process. In this sense, the family of beta processes (1.7) derives from the scale-invariant Poisson process, by taking the "closure under sampling".

Now consider the stable case. Again by Proposition 2.2, ξ_n° is given by the Lévy measure

$$\lambda_{n,a}(s) = ca^{-\alpha}s^{-1-\alpha}(1-s)^n . \tag{3.9}$$

To obtain the full law of $\xi^\circ|\Pi_{1:n}$ additionally requires the law of $\Delta^\circ|\Pi_{1:n}$. To simplify expressions, we consider only the special case $c = \alpha$, though this restriction is not vital. The (surrogate) largest jump $\Delta^\circ \sim P^\circ$ repeatedly appears below raised to a negative power, and it is convenient to express results for Δ° in terms of an equivalent random variable ζ.

COROLLARY 3.3. *Let* $\lambda(s) = \alpha s^{-1-\alpha}$ *and* $\xi^\circ \sim \mathrm{GD}(\lambda, P^\circ)$, *and define* $\zeta := (\Delta^\circ)^{-\alpha}$. *Then*

$$\zeta|\Pi_{1:n} \stackrel{\mathrm{d}}{=} (a^\alpha \psi_n(a) + 1)Y \qquad \textit{where } Y \sim \mathrm{Gamma}(K_n + 1, 1) , \tag{3.10}$$

and $\psi_n(a)$ *is defined as in* (2.4).

We can then marginalize out ξ and determine the predictive distribution [see also 38]:

COROLLARY 3.4. *Choose $\lambda(s) = \alpha s^{-1-\alpha}$, and suppose samples are generated from a Bernoulli process $\mathrm{BeP}(\xi^\circ)$ parametrized by a random measure $\xi^\circ \sim \mathrm{GD}(\lambda, P^\circ)$. Define a random variable ζ_n by $\zeta_n^{-1/\alpha} \sim \mathcal{L}(\Delta^\circ|\Pi_{1:n})$. Then*

$$\Pi_{n+1}|\Pi_1, \ldots, \Pi_n \stackrel{\mathrm{d}}{=} \sum_{j=1}^{N(C_n)} \delta_{U_j} + \sum_{k=1}^{K_n} \mathbf{Z}_{n+1,k}\delta_{U_k} \ , \tag{3.11}$$

where $N(C_n)$ is a $\mathrm{Poisson}(C_n)$ variable, and the variables C_n and $\mathbf{Z}_{n+1,k}$ are conditionally independent given Δ°. Their laws are

$$C_n \stackrel{\mathrm{d}}{=} \zeta_n \alpha B(1-\alpha, n+1) \quad \textit{and} \quad \mathbf{Z}_{n+1,k} \sim \textit{Bernoulli}(J_k) \quad \textit{with} \quad J_k \sim \mathrm{Beta}(n_k - \alpha, n - n_k + 1) \ ,$$

where B is the beta function. The initial sample is marginally distributed as $\Pi_1 \stackrel{\mathrm{d}}{=} \sum_{j=1}^{N(C_0)} \delta_{U_k}$, with $N(C_0) \sim \mathrm{Poisson}(\frac{\zeta\alpha}{1-\alpha})$.

Finally, we can choose a specific surrogate variable Δ°, or equivalently, a specific distribution for ζ. Of particular interest are cases in which the resulting distribution is heavy-tailed [38, 7]. A heavy-tailed law can be generated for example as follows:

COROLLARY 3.5. *Choose ξ° as in Corollary 3.3, and let ζ be an α-stable random variable with density f_ζ. Then*

$$\mathbb{P}(\zeta \in dy|\Pi_{1:n}) \propto y^{K_n} f_\zeta(y) e^{-y\phi_{n,\alpha}} dy \ , \tag{3.12}$$

which is the law of a generalized gamma variable which has been size-biased K_n times.

The law (3.12) is easy to sample, since one can show that a size-biased generalized gamma variable is a mixture of a single generalized gamma variable and gamma variables. It does not yield a closed-form urn scheme, however, unlike the alternative construction explained in Section 5.

3.4. *Constructing the stable-beta by scaling and thinning.* If λ is a stable Lévy density, the random measure $\xi^\circ \sim \mathrm{GD}(\lambda)$ is, conditionally on $\Delta_1 = a$, given by the Lévy density $\lambda_a(s) = a^{-\alpha} c s^{-\alpha-1}$, which partly matches the stable-beta process (1.8), up to a term of the form $(1-s)^\nu$. The actual stable-beta can hence be obtained using conditional thinning, by choosing $h(s) := (1-s)^{\gamma+\alpha-1}$ in Lemma 2.4: If

$$\xi^\circ = \sum_{k\in\mathbb{N}} J_k \delta_{U_k} = \sum_{k\in\mathbb{N}} \frac{\Delta_{k+1}}{\Delta_1} \qquad \text{and} \qquad B_k \sim \mathrm{Bernoulli}((1-J_k)^{\gamma+\alpha-1}) \ , \tag{3.13}$$

for $\gamma > 1-\alpha$, then $\sum_{k\in\mathbb{N}} B_k J_k \delta_{U_k}$ is a stable-beta CRM of the form (1.8).

Since the function h must take values in $[0,1]$ for the construction to be valid, it does not cover the parameter range $\theta \in (-\alpha, 1-\alpha]$. We note the stable-beta can alternatively obtained from a stable using a scaling operation, for a wider range of parameters: The transformation $s \mapsto s/(s+\tau)$ turns $\alpha s^{-1-\alpha}$ into $\alpha\tau^{-\alpha} s^{-1-\alpha}(1-s)^{\alpha-1}$; thinning with $(1-s)^\theta$ then yields a stable-beta.

4. Coupling scaled and normalized random measures. Let $\xi^\circ \sim \mathrm{GD}(\lambda, P^\circ)$ be a random measure constructed as in (1.10). In the following, we generically denote the total mass of any random measure μ on a space $\mathbb{U}$ as $T_\mu := \mu(\mathbb{U})$. Define a random probability measure $\overline{\xi^\circ} := \xi^\circ / T_{\xi^\circ}$. It is obvious to ask whether we can characterize the class of random probability measures so obtained from the class GD; such a characterization is given by Theorem 4.1, which relates GD to the class of Poisson-Kingman measures.

4.1. *Poisson-Kingman measures.* We briefly recall the Poisson-Kingman class [24, 35]: Let μ be a random measure satisfying $T_\mu < \infty$ a.s., so that the random probability measure $\overline{\mu} := \mu / T_\mu$ is well-defined. If μ is in particular a homogeneous CRM with Lévy density λ, $\overline{\mu}$ is called a Poisson-Kingman measure with parameter λ, and its law is denoted $\mathrm{PK}(\lambda)$.

In general, the random variables T_μ and $\overline{\mu}$ are stochastically dependent—the total mass before normalization can carry significant information on the distribution of weights in a given realization of $\overline{\mu}$. The only exception is the Dirichlet process: If a μ is a homogeneous CRM, then $\overline{\mu}$ and T_μ are independent if and only if μ is a gamma process, and $\overline{\mu}$ is hence Dirichlet. If T_μ and $\overline{\mu}$ are dependent, conditioning the law $\mathrm{PK}(\lambda)$ on $T_\mu = t$ yields a family of normalized random measures additionally parametrized by t, denoted

$$\mathrm{PK}(\lambda|t) := \mathcal{L}(\overline{\mu}|T_\mu = t) \; . \tag{4.1}$$

We may now randomize t, by choosing a random variable T° with law Q° on $\mathbb{R}_+$, and define the random probability measure

$$\overline{\mu}^\circ := \overline{\mu}|(T_\mu = T^\circ) \qquad \text{with law} \qquad \mathrm{PK}(\lambda, Q^\circ) := \mathcal{L}(\overline{\mu}^\circ) = \int_{\mathbb{R}_+} \mathrm{PK}(\lambda|t) Q^\circ(dt) \; . \tag{4.2}$$

This is the general class of **Poisson-Kingman measures** in the sense of Pitman [35]. If λ is the Lévy density of a gamma process, $\mathrm{PK}(\lambda)$ is a Dirichlet process. If λ is an α-stable Lévy density and Q° the law of a polynomially tilted stable variable, then $\mathrm{PK}(\lambda)$ is the law of normalized stable process, and $\mathrm{PK}(\lambda, Q^\circ)$ that of a Pitman-Yor process [35].

4.2. *Normalizing scaled subordinators.* We begin with a general result, and then obtain important special cases as corollaries. Start with a random measure $\xi \sim \mathrm{GD}(\lambda)$, and posit an arbitrary joint distribution for a surrogate largest jump Δ°, and a surrogate total mass T°, requiring only absolute continuity $\mathcal{L}(\Delta^\circ, T^\circ) \ll \mathcal{L}(\Delta_1, T_\xi)$. In other words, choose any non-negative, measurable function ω satisfying $\mathbb{E}[\omega(\Delta_1, T_\xi)] = 1$, and define Δ° and T° by

$$\mathbb{P}(\Delta^\circ \in da, T^\circ \in dt) = \omega(a,t)\mathbb{P}(\Delta_1 \in da, T_\xi \in dt) \; . \tag{4.3}$$

Then define the random measure

$$\eta \;:=\; \xi|(\Delta_1 = \Delta^\circ, T_\xi = T^\circ) \;=\; \xi^\circ|(T_\xi = T^\circ) \; , \tag{4.4}$$

generalizing the definition of (1.10) by additionally mixing over the total mass. The next theorem is a more detailed version of Theorem 1.2, and generalizes a result on scale-invariant Poisson processes by Arratia, Barbour, and Tavaré [2].

THEOREM 4.1. *Conditionally on $T_\xi = t$ for any $t \in (0,1]$, the random probability measure $\overline{\eta}$ obtained by normalizing the random measure η defined in (4.4) is a Poisson-Kingman measure. More precisely, if (Q_k) are the ranked weights of $\overline{\eta}$, then*

$$\mathcal{L}(Q_{1:\infty}|T^\circ = t) = \mathrm{PK}(\lambda, P_t^\circ) \qquad \text{where} \qquad P_t^\circ := \mathcal{L}(\Delta^\circ t|T^\circ = t)\ . \tag{4.5}$$

Additionally conditioning on $\Delta^\circ = a$ makes the law independent of the choice of P°, and

$$\mathcal{L}(Q_{1:\infty}|T^\circ = t, \Delta^\circ = a) = \mathrm{PK}(\lambda|at) \tag{4.6}$$

holds for any $t \in (0,1]$.

Two special cases of Theorem 4.1 are of particular interest: One is the basic case $\eta \in \mathrm{GD}(\lambda)$, where $\Delta^\circ = \Delta_1$ and $T^\circ = T_\xi$. The other is a choice of P_t° that lets us obtain arbitrary $\mathrm{PK}(\lambda, \gamma)$ measures. For both cases, we can obtain explicit forms for P_t°. We again write f_{T_Δ} for the Lebesgue density of $T_\Delta := \sum_{k=1}^\infty \Delta_k$.

COROLLARY 4.2. *Let $\eta \in \mathrm{GD}(\lambda, P^\circ)$. Then for any $t \in (0,1]$, identity (4.5) holds with*

$$P_t^\circ(dz) = \frac{z\lambda(z/t)f_{T_\Delta}(z)dz}{\int_0^\infty y\lambda(y/t)f_{T_\Delta}(y)dy}\ . \tag{4.7}$$

Now suppose we wish to obtain $\mathrm{PK}(\lambda, \gamma)$, for an arbitrary measure γ. For simplicity, require the density $h := d\gamma/d\mathcal{L}(T_\Delta)$ exists. The measure γ then has Lebesgue density hf_{T_Δ}.

COROLLARY 4.3. *Let $h : [0,\infty) \to [0,\infty)$ be measurable, with $\mathbb{E}[h(T_\Delta)] = 1$. If Δ° and T° are defined by choosing the density in (4.3) as*

$$\omega(a,t) := \frac{h(at)}{a\lambda(a)} \qquad \text{then} \qquad P_t^\circ(dz) = h(z)f_{T_\Delta}(z)dz \tag{4.8}$$

holds for all $t \in (0,1]$. Any Poisson-Kingman model can be obtained in this manner.

We emphasize that, even though P_t° is now independent of t, it is still necessary to condition on a value $t \le 1$. The requirement $\gamma \ll \mathcal{L}(T_\Delta)$ keeps expressions simple, but a closer look at Theorem 4.1 shows it is not essential; for example, P_t° could be a point mass.

4.3. *Examples.* As examples of Theorem 4.1, we consider the cases where λ is a gamma, scale-invariant, or stable subordinator. The scale-invariant case clarifies how the IBP relates to the Chinese restaurant process. An additional example shows which feature model has an analogous relationship to the two-parameter Chinese restaurant process.

First suppose $\Delta_{1:\infty}$ are the ranked jumps of a gamma process with parameter θ, which has Lévy density

$$\lambda(s) := \theta s^{-1}e^{-s}\ . \tag{4.9}$$

In this case, the Lévy measure λ_a defined in (1.12) for the scaled subordinator is still of gamma type—albeit with jumps bounded by 1—so it is not very surprising that $\overline{\xi}_a$ is a Dirichlet process:

PROPOSITION 4.4 (The gamma case). *Let $\xi \in \mathrm{GD}(\lambda)$, where λ is the gamma Lévy density (4.9). Then conditionally on $T_\xi = t \leq 1$, the normalized measure $\overline{\xi}$ is a Dirichlet process with concentration θ, i.e. its weights have law*

$$\mathcal{L}(Q_{1:\infty}|T_\xi = t) = \mathrm{PD}(0,\theta) \qquad \text{for any } t \in (0,1] \,. \tag{4.10}$$

Although the gamma process yields a one-parameter Poisson-Dirichlet, and hence a CRP, it is not suitable for constructing the IBP. Rather, the CRP and IBP can be constructed jointly from a scale-invariant Poisson [see 2, Theorem 3.1]. In this case, T_ξ has the generalized Dickman distribution $\mathcal{L}(D_\theta)$ described in Section 3.1.

PROPOSITION 4.5 (The IBP and CRP derived from a single random measure). *If $\xi \sim \mathrm{GD}(\lambda)$ for the scale-invariant Poisson density $\lambda(s) = \theta s^{-1}$,*

$$\mathcal{L}(Q_{1:\infty}|T_\xi = t) = \mathrm{PD}(0,\theta) \qquad \text{for all } t \in (0,1] \,. \tag{4.11}$$

The random partition induced by $\overline{\xi}|(T_\xi \leq 1)$ is hence the Chinese restaurant process $\mathrm{CRP}(\theta)$, and the feature model induced by ξ is the $\mathrm{IBP}(\theta)$.

Using a stable subordinator, we can define feature models with Pitman-Yor-like properties. This requires the general class $\mathrm{GD}(\lambda, P^\circ)$, though; the case $\mathrm{GD}(\lambda)$ for stable λ is mostly of theoretical interest:

PROPOSITION 4.6 (The stable case). *If $\xi \sim \mathrm{GD}(\lambda)$ for a subordinator with $\lambda(s) = cs^{-1-\alpha}$,*

$$\mathcal{L}(Q_{1:\infty}|T_\xi = t) = \mathrm{PD}(\alpha,\alpha) \qquad \text{for all } t \in (0,1] \,, \tag{4.12}$$

and $\overline{\xi}|T_\xi = t$ is hence a Pitman-Yor process with parameters (α,α).

That the Dirichlet process can be obtained from the unitary random measure underlying the IBP raises the question which unitary random measure has a similar relationship to the Pitman-Yor process. Obtaining the Pitman-Yor process requires the general formulation in Corollary 4.3.

PROPOSITION 4.7. *Choose an α-stable subordinator, $\lambda(s) = cs^{-1-\alpha}$, and choose ω as in Corollary 4.3, with $h(s) := s^{-\theta}$. Then conditionally on $T_\xi \leq 1$, the random probability measure $\overline{\eta}$ is a Pitman-Yor process with parameters (α,θ), i.e.*

$$\mathcal{L}(Q_{1:\infty}|T_\xi = t) = \mathrm{PD}(\alpha,\theta) \qquad \text{for all } t \in (0,1] \,. \tag{4.13}$$

5. Power laws. This section constructs a model that is not a GD model, but does satisfy the desiderata in Section 1, and for which the distribution of K_n is a power law and can be obtained explicitly.

5.1. *Motivation: The two-parameter Poisson-Dirichlet.* One way to construct the two-parameter Poisson-Dirichlet distribution $\mathrm{PD}(\alpha,\theta)$, loosely based on Proposition 21 in [37], is as follows: Let (Δ_k) be the ranked jumps of an α-stable subordinator, and $(\tilde{\Delta}_k)$ their size-biased permutation. Let $Y := E/\tilde{\Delta}_1$, where E is an independent standard exponential variable. The sequence $(\tilde{\Delta}_2, \tilde{\Delta}_3, \dots | Y = y)$ is then distributed as the jumps of a subordinator with Lévy density

$$\lambda_{\alpha,y}(s) = cs^{-\alpha-1}e^{-y^\alpha s}\ , \tag{5.1}$$

which is a generalized gamma process. Let $T := \tilde{\Delta}_2 + \tilde{\Delta}_3 + \dots$ be its total mass, and mix the normalized process $(\tilde{\Delta}_k/T)_{k\geq 2}$ against a gamma distribution: If $\zeta \sim \mathrm{Gamma}(\theta/\alpha, 1)$, for $\theta > 0$, then

$$\Big(\frac{\tilde{\Delta}_2}{T}, \frac{\tilde{\Delta}_3}{T}, \dots \ \Big|\ Y = \zeta^{1/\alpha}\Big) \sim \mathrm{PD}(\alpha,\theta)\ . \tag{5.2}$$

Observe that the variable ζ, when substituted into (5.1), acts as a random scaling factor *on the Lévy density* of the process. That suggests a similar construction for a unitary random measure: Choose a Lévy density λ supported on $[0,1]$ and a positive scalar random variable ζ, and define a homogeneous random measure as

$$(J_k) \sim \mathrm{Subordinator}(\zeta\lambda) \qquad \text{and} \qquad \mu := \sum_{k\in\mathbb{N}} J_k\delta_{U_k}\ . \tag{5.3}$$

As already noted in Section 1, μ is not generally a GD measure, unless λ is stable or scale-invariant.

5.2. *A heavy-tailed case with tractable urn scheme.* For specific choice of ζ, (5.3) yields a tractable model, with a simple urn scheme and a power law for K_n. Consider a Bessel process of dimension $2(1-\sigma)$, for some $\sigma \in (0,1)$, and denote by L_t the length of its excursion above 0 for which the excursion interval contains the time t. Let T be an independent standard exponential variable, and define ζ as the randomization $\zeta := L_T$. This variable is studied extensively by Bertoin, Fujita, Roynette, and Yor [5], who show it has the remarkably simple density

$$\mathbb{P}(\zeta \in dz) = \frac{\sigma}{\Gamma(1-\sigma)}z^{-\sigma-1}(1-e^{-z})dz \qquad \text{for } \sigma \in (0,1)\ , \tag{5.4}$$

and is hence of the form

$$\zeta := \frac{G}{\beta} \qquad \text{where } G \sim \mathrm{Gamma}(1-\sigma,1) \text{ and } \beta \sim \mathrm{Beta}(\sigma,1)\ . \tag{5.5}$$

We refer to the law of this variable as $\mathrm{BFRY}(\sigma)$, following [11]. Clearly, ζ is heavy-tailed, with more mass in the tail as σ decreases.

Let λ be a Lévy density on $[0,1]$, and define a sequence ϕ_k and a function f as

$$\phi_k := \int_0^1 (1-(1-s)^k)\lambda(s)ds \qquad \text{and} \qquad f(x,a,b) := (xa^b + (1-x)(1+a)^b)^{1/b}\ . \tag{5.6}$$

Specifically, for a stable-beta density $\lambda(s) = \alpha s^{-\alpha-1}(1-s)^{\theta+\alpha-1}$, one obtains

$$\phi_k = \frac{\alpha\Gamma(\theta+\alpha+k-1)\Gamma(1-\alpha)}{\Gamma(\theta+k)}\ ,$$

and the following urn scheme. Let $U_1, U_2, \dots \sim_{\mathrm{iid}} \mathrm{Uniform}[0,1]$, and generate $\mathbf{Z}$ as follows:

1. The first row has $H_1 \sim \text{Poisson}(\zeta\phi_1)$ consecutive non-zero entries, where $\zeta \sim \text{BFRY}(\sigma)$.
2. In the $(n+1)$st row, each of the first K_n entries is selected independently with probability $\frac{n_k-\alpha}{\theta+n}$. Additionally, append

$$H_{n+1}|K_n \sim \text{Poisson}(G_{n+1}\phi_{n+1}f(U_n, \textstyle\sum_{j<k}\phi_j, \sigma - K_n)) \tag{5.7}$$

consecutive non-zero entries, where $G_{n+1} \sim \text{Gamma}(1-\sigma+K_n)$.

REMARK 5.1. The only aspect specific to the stable-beta in the urn scheme above is the probability $\frac{n_k-\alpha}{\theta+n}$ in step (2). The urn scheme holds for any Lévy density λ on $[0,1]$, with ϕ_k computed as in (5.6), if the probabilities $\frac{n_k-\alpha}{\theta+n}$ are replaced appropriately. As we show below, however, the choice $\zeta \sim \text{BFRY}(\sigma)$ is essential: It yields a closed form of $\mathcal{L}(K_n)$.

By definition, the variable $K_1 = H_1$ has a mixed Poisson-BFRY distribution. The same turns out to hold generally for all variables K_n, and we write $K \sim \text{Poisson-BFRY}(\sigma, \tau)$ if

$$K|\zeta \sim \text{Poisson}(\tau\zeta) \qquad \text{for} \qquad \zeta \sim \text{BFRY}(\sigma) \qquad \text{and} \qquad \tau > 0\,. \tag{5.8}$$

The main result regarding the urn scheme is the following:

PROPOSITION 5.2. *Let* $\mathbf{Z}$ *be generated by the urn scheme above.*

1. $\mathbf{Z}$ *is equivalently distributed to a random matrix generated from a random measure* μ *according to* (1.5)*, where* μ *is the random measure* (5.3) *and* $\zeta \sim \text{BFRY}(\sigma)$.
2. *For each* $n \in \mathbb{N}$*, the total number* K_n *of features selected in* $\mathbf{Z}$ *has marginal distribution* $K_n \sim \text{Poisson-BFRY}(\sigma, \sum_{j\leq n}\phi_j)$*. In particular,* K_n *exhibits a power law.*

The proof of Proposition 5.2 follows from the properties of the Poisson-BFRY distribution, which are characterized by the following two lemmas. For any $\sigma \in (0,1)$ and $a, b > 0$, define the function

$$p_{\sigma,a,b}(j,k) := \frac{\Gamma(k+j-\sigma)(a-b)^j}{j!\Gamma(-\sigma)} \frac{a^{-\sigma-j} - (1+a)^{\sigma-j}}{b^{\sigma-k} + (1+b)^{\sigma-k}} \tag{5.9}$$

with arguments $j, k \in \mathbb{N} \cup \{0\}$. A straightforward computation shows:

LEMMA 5.3. *The* Poisson-BFRY(σ, τ) *law has mass function* $\mathbb{P}(H = j) = p_{\sigma,\tau,0}(j, 0)$.

Next, consider a sequence of Poisson-BFRY(σ, τ_n) variables, generated with different parameters τ_n, but coupled by a single BFRY variable σ. Partial sums of such variables inherit additivity from the Poisson distribution; the benign properties of the resulting conditionals account for the existence of the closed-form urn scheme above.

LEMMA 5.4. *Let* $\psi_1, \psi_2, \ldots$ *be positive scalars with partial sums* $\tau_n := \sum_{i\leq n}\psi_i$*. Consider random variables* $\zeta \sim \text{BFRY}(\sigma)$ *and* $H_1, H_2, \ldots$ *satisfying*

$$H_n|\zeta \sim \text{Poisson}(\zeta\psi_n) \qquad \textit{and} \qquad H_n \perp\!\!\!\perp_\zeta (H_1, \ldots, H_{n-1})\,, \tag{5.10}$$

Then the following holds:

1. *The partial sum* $K_n = \sum_{i \le n} H_i$ *has law* $K_n \sim \text{Poisson-BFRY}(\sigma, \tau_n)$.
2. *For each* n, $\mathcal{L}(\zeta | H_1 = h_1, \ldots, H_n = h_n) = \mathcal{L}(\zeta | K_n = \sum_{i \le n} h_i)$, *with density*

$$\mathbb{P}(\zeta \in dz | K_n = k) = \frac{z^{k-\sigma-1} e^{-z\tau_n}(1 - e^{-z})}{\Gamma(k-\sigma)(\tau_n^{\sigma-k} - (1+\tau_n)^{\sigma-k})} \, . \tag{5.11}$$

Moreover, the conditional can be simulated as $\zeta | (K_n = k) \stackrel{\text{d}}{=} V_k G$, *for two random variables* $G \sim \text{Gamma}(k + 1 - \sigma, 1)$ *and* V_k *with density*

$$\mathbb{P}(V_k \in dv) = \frac{(k-\sigma) v^{k-\sigma-1}}{\tau_n^{\sigma-k} - (1+\tau_n)^{\sigma-k}} dv \qquad \text{on} \quad [(\tau_n + 1)^{-1}, \tau_n^{-1}] \, . \tag{5.12}$$

3. *The variable* $H_{n+1} | K_n$ *has mass function* $\mathbb{P}(H_{n+1} = j | K_n = k) = p_{\sigma, \tau_{n+1}, \tau_n}(j, k)$.

6. Distributions of row and column sums. We conclude by considering approximations of the row and column sums, measures in terms of the total variation distance d_{TV} between probability measures.

Let $\Sigma_i^{\text{R}} := \sum_k \mathbf{Z}_{ik}$ be the sum of the ith row. Since T_ξ if almost surely finite, so is Σ_i^{R}. If ξ is a homogeneous CRM, $\Sigma_i^{\text{R}} \sim \text{Poisson}(\mathbb{E}[T_\xi])$ by Proposition 2.1. If ξ is not completely random, one may still approximate it by a Poisson variable; the expected error can be related to the first size-biased weight of ξ:

PROPOSITION 6.1. *Let* $\xi \stackrel{\text{d}}{=} \sum_k W_k \delta_{U_k}$ *be a unitary random measure, and* $\mathbf{Z}$ *a random matrix generated from* ξ *as in* (1.5). *Suppose* Σ_i^R *is approximated by a Poisson variable with mean* T_ξ, *and measure approximation error as* $E_\xi := d_{TV}\big(\mathcal{L}(\Sigma_i^R | \xi) - \text{Poisson}(T_\xi)\big)$. *Then* $E_\xi | \xi < \sum_k W_k^2$, *and the expected error is bounded as*

$$\mathbb{E}[E_\xi] < \mathbb{E}[T_\xi \tilde{W}_1] \, , \tag{6.1}$$

where $(\tilde{W}_1, \tilde{W}_2, \ldots)$ *is a size-biased permutation of the weight sequence* (W_k).

Now consider the kth column. Since the column entries are conditionally i.i.d. with mean W_k, the behavior of column sums reduces to that of the weights W_k. The next result provides a tool that relates these weights to those of a random probability measure, along the lines of Theorem 4.1. The bound considers only those weights not exceeding a given size β; all larger weights are discarded from the total mass, by defining

$$\tau_\beta(x) := \sum_{i \in \mathbb{N}} x_i \mathbb{I}\{x_i \le \beta\} \qquad \text{for } \beta > 0 \tag{6.2}$$

for a point set $x = \{x_1, x_2, \ldots\}$ in $\mathbb{R}_+$. Comparing the laws of the remaining jumps using total variation distance then reduces to comparing the scalar variables τ_β to its conditional given the total mass:

PROPOSITION 6.2. *Let* $X = \{X_1, X_2, \ldots\}$ *be the set of jump heights of a subordinator with with Lévy density* λ, *and require that its total mass* $T = \sum X_k$ *a density* f_T *on* $\mathbb{R}_+$.

Let $V = \{V_1, V_2, \ldots\}$ *be the weights of a random probability measure with law* $\mathrm{PK}(\lambda, P_t^\circ)$ *as in* (4.5)*, for some* $t \in (0, 1]$*. Then*

$$d_{TV}(\mathcal{L}(V \cap [0, \beta]), \mathcal{L}(X \cap [0, \beta])) = d_{TV}(\mathcal{L}(\tau_\beta(X)|T = t), \mathcal{L}(\tau_\beta(X))) \tag{6.3}$$

holds for all $\beta \in [0, 1]$*.*

The bound (6.3) generalizes a result for the scale-invariant Poisson process obtained in [2]. As Theorem 4.1, conditioning on $T \leq t = 1$ on the right-hand side of (6.3) is essential; for arbitrary values of T, the result does not generally hold.

Acknowledgments. LFJ was supported in part by grant RGC-HKUST 601712 of the HKSAR. PO was supported in part by grant FA9550-15-1-0074 of AFOSR. YWT's research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 617071.

PROOFS

All results summarized in Section 1 follow from results in later sections, with the exception of Theorem 1.1.

PROOF OF THEOREM 1.1. Following Kingman [24, Appendix 1], we note Δ_1 has density (1.16), so its cdf is $F_{\Delta_1}(s) = e^{-\Lambda(s)}$, and $e^{-\Lambda(\Delta_1)}$ is a uniform variable. Hence, $\Lambda(\Delta_1) \stackrel{\mathrm{d}}{=} E_1$, where E_1 is a unit-rate exponential. Now apply a similar argument conditionally for $k > 1$: The ranked jumps of a subordinator have the Markov property

$$\Delta_k|\Delta_{k-1} = y, \Delta_{k-2}, \ldots, \Delta_1 \quad \stackrel{\mathrm{d}}{=} \quad \Delta_k|\Delta_{k-1} = y \; . \tag{A.4}$$

As $\Delta_{2:\infty}|(\Delta_1 = y)$ is a subordinator with Lévy density $\lambda(s)\mathbb{I}\{s \leq y\}$, the conditional density and cdf of its largest jump Δ_2 are, respectively,

$$f(s|y) = \lambda(s)e^{-(\Lambda(s)-\Lambda(y))} \qquad \text{and} \qquad F(s|y) = e^{-(\Lambda(s)-\Lambda(y))} \qquad \text{for } s \leq y \; . \tag{A.5}$$

Iterating the argument shows the conditional density and cdf of $\Delta_k|\Delta_{k-1} = y$ are again given by (A.5). By (A.4), the variable $\Delta_k|\Delta_{k-1} = y, \Delta_{k-2}, \ldots, \Delta_1$ hence has conditional cdf $F(s, y)$. The variables $M_k := \Lambda^{-1}(E_1 + \ldots + E_k)$ therefore have distribution $(M_k) \stackrel{\mathrm{d}}{=} (\Delta_k)$. Conditioning on $\Delta_1 = a$, and hence on $E_1 = \Lambda(a)$, yields (1.17). □

A.1. Proofs for Section 2. For the proof of Proposition 2.1, recall the Laplace functional of a random measure μ, for a non-negative Borel function g, is defined as $\mathbb{E}[e^{-\mu(g)}]$, and μ is Poisson iff

$$\mathbb{E}[e^{-\mu(g)}] = e^{-(\mathbb{E}[\mu])(1-e^{-g})} \; , \tag{A.6}$$

see e.g. [21, Lemma 12.2].

PROOF OF PROPOSITION 2.1. We show that the random measure Π satisfies (A.6) if μ is a CRM. The definition in (2.1) implies $\mathbb{E}[\Pi] = \mathbb{E}[\mu]$, and since μ is a homogeneous CRM, $\mathbb{E}[\Pi(\bullet)] = \mathbb{E}[\mu(\mathbb{U})]G(\bullet)$ then follows as claimed.

Conditionally on the weights W_k of $\mu = \sum_{k\in\mathbb{N}} W_k\delta_{U_k}$, the variables Z_k in (2.1) are independent Bernoulli, each with expectation W_k, hence

$$\mathbb{E}[e^{-\Pi(g)}|(W_k)] = \prod_k \mathbb{E}[e^{-Z_k g(X_k)}|(W_k)] = \prod_k ((1 - W_k) + W_k\mathbb{E}[e^{-g(X_1)}]) \; . \tag{A.7}$$

If we abbreviate $h(s) := -\log((1 - s) + s\mathbb{E}[e^{-g(X_1)}])$, we obtain

$$\mathbb{E}[e^{-\Pi(g)}|(W_k)] = e^{-\sum_k h(W_k)} \; . \tag{A.8}$$

The exponent on the right can be written as $\sum h(W_k) = N(h)$, where $N := \sum \delta_{W_k}$. The Laplace functional of Π at g then coincides with that of N at h:

$$\mathbb{E}[e^{-\Pi(g)}] = \mathbb{E}[\mathbb{E}[e^{-\Pi(g)}|(W_k)]] = \mathbb{E}[e^{-N(h)}] \tag{A.9}$$

Since (W_k) are the jumps of a homogeneous CRM, N is a Poisson random measure with expectation $\mathbb{E}[N(ds)] = \lambda(s)ds$. By (A.6), its Laplace functional at h is

$$\mathbb{E}[e^{-N(h)}] = \mathbb{E}[e^{-\int(1-e^{-h(s)})\lambda(s)ds}] \; . \tag{A.10}$$

Since

$$1 - e^{-h(s)} = s(1 - \mathbb{E}[e^{-g(U_1)}]) = s \int_{\mathbb{U}} 1 - e^{-g(u)} G(du) \; , \tag{A.11}$$

the integral in the exponent of (A.10) is

$$\int_0^\infty (1 - e^{-h(s)})\lambda(s)ds = \int_0^\infty (s \int_{\mathbb{U}} 1 - e^{-g(u)} G(du))\lambda(s)ds = \mathbb{E}[\mu(\mathbb{U})]G(1 - e^{-g}) \; . \tag{A.12}$$

Hence, $\mathbb{E}[e^{-\Pi(g)}] = \mathbb{E}[e^{-\mathbb{E}[\mu(\mathbb{U})]G(1-e^{-g})}]$, and Π is indeed Poisson by (A.6). □

Proof of Proposition 2.2. By definition, ξ° is completely random given Δ°, which implies (2.3), and the conditional distributions of ξ_n° and the conditional weights J_{nk}^* in (2.5) and (2.6) can be derived from the results of Kim [22]. To obtain the conditional (2.7) of Δ°, we note that all information in a sample $\Pi_1, \ldots, \Pi_n$ relevant to conditioning is summarized by the K_n atom locations U_k^* and their multiplicities n_k. The likelihood of observing the sample given $\Delta^\circ = a$ is

$$\mathbb{P}(du_{1:K_n}^*, n_{1:K_n}, n | \Delta^\circ = a) = e^{-\psi_n(a)} \prod_{k=1}^{K_n} c_a(n, n_k) G(du_k^*) \; . \tag{A.13}$$

If Δ° has density f_{Δ°, Bayes' theorem hence yields (2.7). □

Proof of Lemma 2.4. Let G denote the law of U_1. Since (J_k, B_k, U_k) are the points of a marked Poisson process, the Laplace transform $\mathbb{E}[\mathrm{e}^{-\mu_h(f)}] = \mathrm{e}^{-\psi(f)}$ of μ_h is given by

$$\begin{aligned} \psi(f) &= \sum_{a=0}^{1} \int_{\mathbb{U}} \int_0^\infty (1 - \mathrm{e}^{-af(y)s}) h^a(s)[1 - h(s)]^{1-a} \lambda(s) ds G(du) \\ &= \int_{\mathbb{U}} \int_0^\infty (1 - \mathrm{e}^{-f(y)s}) h(s) \lambda(s) ds G(du) \; , \end{aligned} \tag{A.14}$$

since $(1 - \mathrm{e}^{-af(y)s}) = 0$ for $a = 0$. □

A.2. Proofs for Section 4. This section gives the proof of Theorem 4.1, which we break down into a number of lemmas. Let f_{Δ_1} be the density of Δ_1, and define the function

$$g_\lambda(a, t) = a\lambda(a) f_{\Delta_1}(at) \; . \tag{A.15}$$

We first note the following simple form for the joint density of (Δ_1, T_ξ), which is valid only if T_ξ takes values in $[0, 1]$; outside this range, the expressions become considerably more complicated.

Lemma A.3. *Let $\xi \sim \mathrm{GD}(\lambda)$ for a Lévy density λ. Then $f_{\Delta_1, T_\xi}(a, t) = g_\lambda(a, t)$ for $t \leq 1$.*

Proof. The total mass T_Δ and largest normalized jump V of the subordinator $(\Delta_1, \Delta_2, \ldots)$ are related to the variables Δ_1 and T_Δ by

$$V = \frac{\Delta_1}{T_\Delta} = \frac{1}{1 + T_\xi} \qquad \text{and} \qquad T_\Delta = (1 + T_\xi)\Delta_1 \; . \tag{A.16}$$

The hypothesis $T_\xi \leq 1$ is hence equivalent to $V \geq \frac{1}{2}$. Let f_{T_Δ} be the Lebesgue density of T_Δ. By [37, Proposition 45], the joint density of T_Δ and V is $f_{V,T_\Delta}(a,\tau) = \tau\lambda(\tau a) f_{T_\Delta}(\tau - a\tau)$, provided that $V \geq \frac{1}{2}$. The change of variables $\tau \mapsto (1+t)a$ and $v \mapsto \frac{1}{1+t}$, and renormalizing with respect to t, yields $f_{\Delta_1,T_\xi}(a,t) = \frac{1}{1+t} f_{V,T_\Delta}\big((1+t)a, \frac{1}{1+t}\big)$, which is just $g_\lambda(a,t)$. □

Next, consider the effect of substituting λ by λ_a in a Poisson-Kingman partition.

LEMMA A.4. *Let λ be a Lévy measure and $\lambda_a(s) = a\lambda(as)\mathbb{I}\{s \leq 1\}$, for some $a > 0$. Then $\mathrm{PK}(\lambda_a|t) = \mathrm{PK}(\lambda|at)$ whenever $t \in [0,1]$.*

PROOF. By [35, Theorem 4], the EPPF of a $\mathrm{PK}(\nu|t)$ partition for some Lévy density ν is given by

$$p_\nu\big(|A_1|,\ldots,|A_k|\big|t\big) = \int_0^t \frac{f(t-s)}{t^n f(t)} s^{n+K-1} \int_{\mathcal{S}_k} \Big(\prod_{k=1}^K \nu(su_k) u_k^{|A_k|}\Big) du_{1:K-1} ds \,, \tag{A.17}$$

where $\mathcal{S}_K$ denotes the standard simplex in $\mathbb{R}^K$, $du_{1:K-1} = du_1 \cdots du_{K-1}$, and f is the density of the total mass of the random measure defined by ν. Hence,

$$p_{\lambda_a}\big(|A_1|,\ldots,|A_k|\big|t\big) = \int_0^t \frac{f_\Delta(t-s)}{t^n f_\Delta(t)} \frac{s^{n+K-1}}{a^K} \int_{\mathcal{S}_k} \Big(\prod_{k=1}^K \lambda(asu_k)\mathbb{I}\{su \leq 1\} u_k^{|A_k|}\Big) du_{1:K-1} ds \,.$$

Since $s \leq t \leq 1$ by hypothesis, and $u_k \leq 1$, the indicator $\mathbb{I}\{s\}u \leq 1$ is always 1, and a change of variables $v(s) := s/a$ yields

$$p_{\lambda_a}\big(|A_1|,\ldots,|A_k|\big|t\big) = \int_0^{at} \frac{f_\Delta(at-s)}{a^n t^n f_\Delta(at)} s^{n+K-1} \int_{\mathcal{S}_k} \Big(\prod_{k=1}^K \lambda(su_k) u_k^{|A_k|}\Big) du_{1:K-1} ds \,,$$

which by (A.17) is just $p_\lambda\big(|A_1|,\ldots,|A_k|\big|at\big)$. □

PROOF OF THEOREM 4.1. Since $T_\xi \leq 1$, the conditional density of T_ξ given $\Delta_1 = a$ is

$$f_{T_\xi|\Delta_1}(t|a) = \frac{g_\lambda(a,t)}{\int_0^1 g_\lambda(a,s)ds} = \frac{a\lambda(as) f_{T_\Delta}(t)}{\int_0^1 a\lambda(as) f_{T_\Delta}(s)ds} \qquad \text{for } t \in [0,1], \tag{A.18}$$

by Lemma A.3. The random measure $\overline{\xi}|(\Delta_1 = a)$ is hence a Poisson-Kingman measure with Lévy density λ_a, and by Lemma A.4, we have

$$\mathcal{L}(\overline{\xi}|\Delta_1 = a, T_\xi = t) = \mathrm{PK}(\lambda_a|t) = \mathrm{PK}(\lambda|at) \qquad \text{whenever } t \leq 1 \,, \tag{A.19}$$

which proves (4.6). To establish (4.5), note the conditional density of $\Delta^\circ|T^\circ$ satisfies $f_{\Delta^\circ|T^\circ}(a|t) \propto \omega(a,t)g(a,t)$ by Lemma A.3. Hence,

$$\mathcal{L}(\overline{\xi}|T^\circ = t) = \int_{\mathbb{R}_+} \mathcal{L}(\overline{\xi}|\Delta^\circ = a, T^\circ = t) f_{\Delta^\circ|T^\circ}(a|t)da \propto \int_{\mathbb{R}_+} \mathrm{PK}(\lambda|at)\omega(a,t)g(a,t)da \,.$$

Changing variables to $z(a) = at$ yields

$$\mathcal{L}(\overline{\xi}|T^\circ = t) \propto \int_{\mathbb{R}_+} \mathrm{PK}(\lambda|z)\omega(\tfrac{z}{t},t)g(\tfrac{z}{t},t)tdz \propto \int_{\mathbb{R}_+} \mathrm{PK}(\lambda|z)\tfrac{1}{t}\omega(\tfrac{z}{t},t)g(\tfrac{z}{t},t)dz\ . \tag{A.20}$$

Since $\mathrm{PK}(\lambda, P_t^\circ) = \int_{\mathbb{R}_+} \mathrm{PK}(\lambda|z)P_t^\circ(dz)$, indeed $P_t^\circ = \mathcal{L}(\Delta^\circ t|T^\circ = t)$ as claimed. □

The examples in Proposition 4.4–Proposition 4.7 are obtained as follows: In the gamma case, in Proposition 4.4, the total mass T_Δ of the subordinator has distribution $\mathrm{Gamma}(\theta, 1)$. For this choice of λ, Poisson-Kingman model satisfy $\mathrm{PK}(\lambda|t) = \mathrm{PK}(\lambda) = \mathrm{PD}(0,\theta)$ for all $t > 0$ (see [35, §5.1]). Substituting into Corollary 4.2 hence yields

$$\mathcal{L}(Q_{1:\infty}|T_\xi = t) = \mathrm{PK}(\lambda, P_t^\circ) = \mathrm{PK}(\lambda) = \mathrm{PD}(0,\theta) \qquad \text{for } t \in (0,1]\ . \tag{A.21}$$

Note that, in this case, $P_t^\circ(dz) \propto z^{\theta-1}e^{-z(1+t)/t}$.

In the scale-invariant case $\lambda(s) = \theta s^{-1}$ in Proposition 4.5, $\xi|(\Delta_1 = a)$ has Lévy density $\lambda_a = \theta s^{-1}\mathbb{I}\{s \leq 1\}$ by (1.12), and the total mass T_ξ is a Dickman variable D_θ as in Section 3.1. Theorem 4.1 shows $\mathcal{L}(Q_{1:\infty}|T_\xi = t, \Delta^\circ = a) = \mathrm{PK}(\lambda|D_\theta = at)$. To evaluate $\mathrm{PK}(\lambda|D_\theta = at)$, we reduce to a gamma Lévy density $\lambda'(s) = \theta s^{-1}e^{-s}$, using the invariance

$$\mathrm{PK}(e^{-as}\lambda|t) = \mathrm{PK}(\lambda|t) \qquad \text{for all } t > 0 \tag{A.22}$$

of PK measures under exponential tilting [35, §4.2]. Since exponentially tilting λ_a yields

$$e^{-as}\lambda_a(s) = \theta s^{-1}e^{-as}\mathbb{I}\{s \leq 1\} = \lambda'_a(s)\ , \tag{A.23}$$

we obtain $\mathrm{PK}(\lambda|D_\theta = at) = \mathrm{PK}(\lambda'|T_\xi = at) = \mathrm{PD}(0,\theta)$.

In Proposition 4.6, the subordinator is α-stable, hence $\lambda(s) = cs^{-1-\alpha}$ and $T_\Delta \sim \mathrm{Stable}(\alpha)$. By Corollary 4.2, $P_t^\circ(dz) \propto z^{-\alpha-1}f_{T_\Delta}(z)dz$, which is the law of a polynomially α-tilted, α-stable variable $S_{\alpha,\alpha}$. Hence,

$$\mathcal{L}(Q_{1:\infty}|T_\xi = t) = \mathrm{PK}(\lambda, \mathcal{L}(S_{\alpha,\alpha}))\ , \tag{A.24}$$

which is the $\mathrm{PD}(\alpha,\alpha)$ distribution, see [35].

In Proposition 4.7, λ is an α-stable subordinator and T_Δ is hence an α-stable random variable with density f_α. By Corollary 4.3,

$$\mathcal{L}(Q_{1:\infty}|T_\xi = t) = \mathrm{PK}(\lambda, z^{-\theta}f_\alpha(z))\ , \tag{A.25}$$

which is the $\mathrm{PD}(\alpha,\theta)$ distribution [35].

A.3. Proofs for Section 6. The rows sums follow a generalized binomial distribution, the Poisson approximation error can be bounded using Le Cam's inequality [25]:

Proof of Proposition 6.1. By Le Cam's inequality, the law of Σ_i^{R} is approximately Poisson, with approximation error $\sum_k W_k^2$. Since $T_\xi < \infty$ a.s. by hypothesis, ξ can be normalized, and we write $\tilde{W}_1$ for the first size-biased pick from $W_{1:\infty}$. The variable $\tilde{V}_1 := \tilde{W}_1/T_\xi$

is hence the first size-biased weight of a random probability measure, and satisfies the identity $\mathbb{E}[g(\tilde{V}_1)] = \mathbb{E}[\sum_k \tilde{V}_k g(\tilde{V}_k)]$ for any measurable $g : (0,\infty) \to (0,\infty)$, see [36, Eq. (2.22)]. Choosing $g(v) = t^2 v$,

$$\mathbb{E}\Big[\sum_k W_k^2\Big] = \mathbb{E}_{T_\xi}\Big[\mathbb{E}\Big[\sum_k t^2 \tilde{V}_k^2 \Big| T_\xi = t\Big]\Big] = \mathbb{E}_{T_\xi}\big[\mathbb{E}[t^2 \tilde{V}_1 | T_\xi = t]\big] = \mathbb{E}[T_\xi^2 \tilde{V}_1] \tag{A.26}$$

as claimed. □

The next proof uses the following auxiliary result, which paraphrases [2, Corollary 4.1]:

LEMMA A.5. *Let P and P' be two probability measures on $\mathcal{X}$ with $P \ll P'$, and let $\phi : \mathcal{X} \to \mathcal{Y}$ be a measurable mapping into a space $\mathcal{Y}$. Suppose ϕ is sufficient for P and P'; that is, there is some function f such that $dP/dP'(x) = f(\phi(x))$. Then the total variation distance between the image measures $\phi(P)$ and $\phi(P')$ satisfies $d_{TV}(P, P') = d_{TV}(\phi(P), \phi(P'))$.*

PROOF OF PROPOSITION 6.2. Abbreviate $X_\beta := X \cap [0, \beta]$. From Theorem 4.1, we may substitute $\mathcal{L}(X_\beta | T = t)$ for $\mathcal{L}(V)$ in (6.3). By [2, Corollary 4.1], we hence have to show that τ_β is a sufficient statistic for the pair $\mathcal{L}(X_\beta)$ and $\mathcal{L}(X_\beta | T = t)$. Let L be Lebesgue measure on $\mathbb{R}_+$. Since T has a density $f_T = d\mathcal{L}(T)/dL$, the conditional density $f_{T|X_\beta} = d\mathcal{L}(T|X_\beta)/dL$ exists. As the continuous-time process $(T_\beta)_{\beta \in \mathbb{R}_+}$ is Markovian, T and X_β are conditionally independent given T_β, and $f_{T|X_\beta}$ decomposes as $f_{T|X_\beta}(t|x) = \tilde{f}(t|\tau_\beta(x))$ for a suitable conditional density function $\tilde{f}$. The joint law then has density

$$\frac{\mathbb{P}(X_\beta \in dx, T \in dt)}{\mathbb{P}(X_\beta \in dx) \otimes \lambda(dt)} = f_{T|X_\beta}(t|x)\mathbf{1}(x) \,, \tag{A.27}$$

where $\mathbf{1}$ is the constant function with value 1. This means the density

$$\frac{\mathbb{P}(X_\beta \in dx | T = t)}{\mathbb{P}(X_\beta \in dx)} = \frac{f_{T|X_\beta}(t|x)}{f_T(t)} = \frac{\tilde{f}(t|\tau_\beta(x))}{f_T(t)} \,. \tag{A.28}$$

depends on X only through τ_β, and τ_β is indeed sufficient for $(\mathcal{L}(X|T), \mathcal{L}(X))$. □

REFERENCES

[1] Arratia, R. (1998). On the central role of scale invariant Poisson processes on $(0, \infty)$. In *Microsurveys in discrete probability (Princeton, NJ, 1997)*, volume 41 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 21–41. Amer. Math. Soc., Providence, RI.

[2] Arratia, R., Barbour, A. D., and Tavaré, S. (1999). The Poisson-Dirichlet distribution and the scale-invariant Poisson process. *Combin. Probab. Comput.*, **8**(5), 407–416.

[3] Arratia, R., Barbour, A. D., and Tavaré, S. (2003). *Logarithmic combinatorial structures: A probabilistic approach*. EMS Monographs in Mathematics. European Mathematical Society.

[4] Berti, P., Crimaldi, I., Pratelli, L., and Rigo, P. (2015). Central limit theorems for an Indian buffet model with random weights. *Ann. Appl. Probab.*, **25**(2), 523–547.

[5] Bertoin, J., Fujita, T., Roynette, B., and Yor, M. (2006). On a particular class of self-decomposable random variables: the durations of Bessel excursions straddling independent exponential times. *Probab. Math. Statist.*, **26**(2), 315–366.

[6] Broderick, T., Jordan, M. I., and Pitman, J. (2011). Beta processes, stick-breaking, and power laws.

[7] Broderick, T., Jordan, M. I., and Pitman, J. (2012). Beta processes, stick-breaking and power laws. *Bayesian Anal.*, **7**(2), 439–475.

[8] Broderick, T., Jordan, M. I., and Pitman, J. (2013). Cluster and feature modeling from combinatorial stochastic processes. *Statist. Sci.*, **28**(3), 289–312.

[9] Caron, F. (2012). Bayesian nonparametric models for bipartite graphs. In *Advances in Neural Information Processing Systems 25*, pages 2051–2059.

[10] Daley, D. and Vere-Jones, D. (2008). *An introduction to the theory of point processes*, volume I and II. Springer, 2nd edition.

[11] Devroye, L. and James, L. (2014). On simulation and properties of the stable law. *Stat. Methods Appl.*, **23**(3), 307–343.

[12] Ferguson, T. S. and Phadia, E. G. (1979). Bayesian nonparametric estimation based on censored data. *Annals of Statistics*, **7**, 163–186.

[13] Fox, E., Hughes, M., Sudderth, E., and Jordan, M. (2014). Joint modeling of multiple related time series via the beta process with application to motion capture segmentation. *Annals of Applied Statistics*, **8**(3), 1281–1313.

[14] Ghahramani, Z., Griffiths, T. L., and Sollich, P. (2006). Bayesian nonparametric latent feature models. In *Bayesian Statistics 8 (Valencia, 2006)*. Oxford Univ. Press.

[15] Griffiths, T. L. and Ghahramani, Z. (2006). Infinite latent feature models and the Indian buffet process. In *Adv. Neural Inf. Process. Syst.*

[16] Griffiths, T. L. and Ghahramani, Z. (2011). The Indian buffet process: An introduction and review. *J. Mach. Learn. Res.*, **12**, 1185–1224.

[17] Heaukulani, C. and Roy, D. M. (2015). The combinatorial structure of beta negative binomial processes. *Bernoulli*, to appear.

[18] Hjort, N. L. (1990). Nonparametric Bayes estimators based on beta processes in models for life history data. *Ann. Statist.*, **18**, 1259–1294.

[19] Jacob, L., Vert, J.-P., and Bach, F. R. (2009). Clustered multi-task learning: A convex formulation. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 745–752. Curran Associates, Inc.

[20] Janson, R. G. S., Luczak, T., and Rucinski, A. (2000). Random graphs. J. Wiley & Sons.

[21] Kallenberg, O. (2001). *Foundations of Modern Probability*. Springer, 2nd edition.

[22] Kim, Y. (1999). Nonparametric Bayesian estimators for counting processes. *Ann. Statist.*, **27**(2), 562–588.

[23] Kingman, J. F. C. (1967). Completely random measures. *Pacific Journal of Mathematics*, **21**(1), 59–78.

[24] Kingman, J. F. C. (1975). Random discrete distributions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **37**, 1–22.

[25] Le Cam, L. (1960). An approximation theorem for the Poisson binomial distribution. *Pacific J. Math.*, **10**, 1181–1197.

[26] Lijoi, A. and Prünster, I. (2009). Distributional properties of means of random probability measures. *Stat. Surv.*, **3**, 47–95.

[27] Lounici, K., Tsybakov, A. B., Pontil, M., and van de Geer, S. A. (2009). Taking advantage of sparsity in multi-task learning. In *Conference on Learning Theory*, volume 1050, page 9.

[28] Meeds, E., Ghahramani, Z., Neal, R. M., and Roweis, S. T. (2007). Modeling dyadic data with binary

latent factors. In *Advances in Neural Information Processing Systems 19*, pages 977–984.
[29] Miller, K., Jordan, M. I., and Griffiths, T. L. (2009). Nonparametric latent feature models for link prediction. In *Advances in neural information processing systems*, pages 1276–1284.
[30] Navarro, D. J. and Griffiths, T. L. (2008). Latent features in similarity judgments: A nonparametric bayesian approach. *Neural computation*, **20**(11), 2597–2628.
[31] Ni, K., Paisley, J., Carin, L., and Dunson, D. (2008). Multi-task learning for analyzing and sorting large databases of sequential data. *Signal Processing, IEEE Transactions on*, **56**(8), 3918–3931.
[32] Paisley, J., Zaas, A., Ginsburg, G., Woods, C., and Carin, L. (2010). A stick-breaking construction of the beta process. In *ICML*.
[33] Penrose, M. D. and Wade, A. R. (2004). Random minimal directed spanning trees and Dickman-type distributions. *Adv. in Appl. Probab.*, **36**(3), 691–714.
[34] Perman, M. (1993). Order statistics for jumps of normalised subordinators. *Stochastic Process. Appl.*, **46**(2), 267–281.
[35] Pitman, J. (2003). Poisson-Kingman partitions. In *Statistics and science: a Festschrift for Terry Speed*, volume 40 of *IMS Lecture Notes Monogr. Ser.*, pages 1–34. Inst. Math. Statist., Beachwood, OH.
[36] Pitman, J. (2006). *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Springer.
[37] Pitman, J. and Yor, M. (1997). The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann. Probab.*, **25**(2), 855–900.
[38] Teh, Y.-W. and Görur, D. (2009). Indian buffet processes with power-law behavior. In *Adv. Neural Inf. Process. Syst.*
[39] Teh, Y.-W., Görur, D., and Ghahramani, Z. (2007). Stick-breaking construction for the Indian buffet process. In *AISTATS*.
[40] Thibaux, R. and Jordan, M. I. (2007). Hierarchical beta processes and the Indian buffet process. In *J. Mach. Learn. Res. Proceedings*, volume 2, pages 564–571.
[41] Titsias, M. K. (2008). The infinite gamma-poisson feature model. In *Advances in Neural Information Processing Systems 20*, pages 1513–1520.
[42] Zhou, M., Hannah, L. A., Dunson, D. B., and Carin, L. (2012). Beta-negative binomial process and Poisson factor analysis. In *Proceedings of AISTATS 2012 (JMLR W&CP 22)*, pages 1462–1471.

Deparment of Information Systems, Business
Statistics, and Operations Management
Clear Water Bay, Kowloon
Hong Kong
E-mail: lancelot@ust.hk

Department of Statistics
1255 Amsterdam Avenue
New York, NY-10027, USA
E-mail: porbanz@stat.columbia.edu

Department of Statistics
1 South Parks Road
Oxford OX1 3TG, U.K.
E-mail: y.w.teh@stats.ox.ac.uk